

Old Dominion University

**ODU Digital Commons**

---

Engineering Management & Systems  
Engineering Faculty Publications

Engineering Management & Systems  
Engineering

---

2015

## **Research Agenda into Human-Intelligence/Machine-Intelligence Governance**

Teddy Steven Cotter

Follow this and additional works at: [https://digitalcommons.odu.edu/emse\\_fac\\_pubs](https://digitalcommons.odu.edu/emse_fac_pubs)



Part of the [Artificial Intelligence and Robotics Commons](#), [Databases and Information Systems Commons](#), [Software Engineering Commons](#), and the [Theory and Algorithms Commons](#)

---

## **RESEARCH AGENDA INTO HUMAN-INTELLIGENCE/MACHINE-INTELLIGENCE GOVERNANCE**

**T. Steven Cotter**  
**Old Dominion University**  
**tcotter@odu.edu**

---

### **Abstract**

Since the birth of modern artificial intelligence (AI) at the 1956 Dartmouth Conference, the AI community has pursued modeling and coding of human intelligence into AI reasoning processes (HI  $\Rightarrow$  MI). The Dartmouth Conference's fundamental assertion was that every aspect of human learning and intelligence could be so precisely described that it could be simulated in AI. With the exception of knowledge specific areas (such as IBM's Big Blue and a few others), sixty years later the AI community is not close to coding global human intelligence into AI. In parallel, the knowledge management (KM) community has pursued understanding of organizational knowledge creation, transfer, and management (HI  $\Rightarrow$  HI) over the last 40 years. Knowledge management evolved into an organized discipline in the early 1990's through formal university courses and creation of the first chief knowledge officer organizational positions. Correspondingly, over the last 25 years there has been growing research into the transfer of intelligence and cooperation among computing systems and automated machines (MI  $\Rightarrow$  MI). In stark contrast to the AI community effort, there has been little research into transferring AI knowledge and machine intelligence into human intelligence (MI  $\Rightarrow$  HI) with a goal of improving human decision making. Most important, there has been no research into human-intelligence/machine-intelligence decision governance; that is, the policies and processes governing human-machine decision making toward systemic mission accomplishment. To address this gap, this paper reports on a research initiative and framework toward developing an HI-MI decision governance body of knowledge and discipline.

### **Keywords**

Governance, Human Intelligence, Machine Intelligence.

### **Introduction**

With the assistance of increasing computing power, human knowledge about human knowledge and intelligence grew rapidly in the last half of the 20<sup>th</sup> century and continues to accelerate in the 21<sup>st</sup> century. Increasing computing power facilitated the growth in artificial intelligence, which has an explicit goal of encoding human intelligence into computer artificial intelligence. On the human side, much of human knowledge about human knowledge is encoded and managed in domain-specific knowledge bases. However, to date the goals of capturing human tacit knowledge and converting it into explicit, actionable knowledge and of achieving true artificial intelligence are far from being achieved. In the first of a series of stories on the Public Broadcasting System about the future of AI, Hari Sreenivasin (2015) noted that AI enabled "... self-driving cars have been test-driven, without incident, for hundreds of thousands of miles, but are not quite ready for consumers." His interviewee, Fei-Fei Li, Director of the Stanford Artificial Intelligence Laboratory, noted that, "Yes, we have prototype cars that can drive by themselves. But without smart vision, they cannot really tell the difference between a crumpled paper bag, which can be run over, and a rock that size, which should be avoided." Similarly, in his book *The Digital Doctor*, Robert Wachter (2015, p. xiii) notes that capturing medical knowledge and integrating it into practice "... turns out to have been magical thinking ... we're learning that computers make some things better, some things worse, and they change everything."

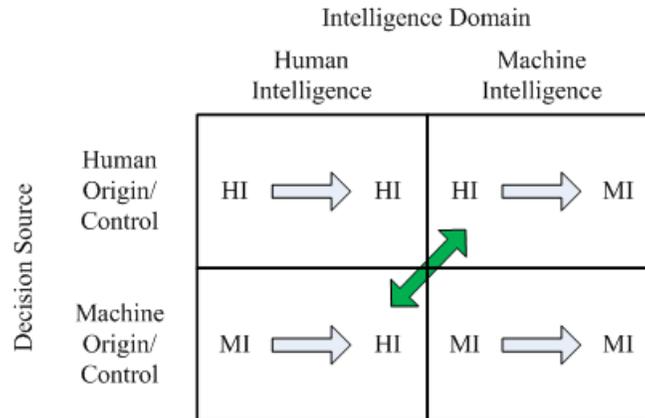
Regardless of the difficulties, research and movement toward integrating artificial intelligence into human organizations and everyday life continues to accelerate. However, only recently has research been initiated into the management of the human-intelligent/machine-intelligent organization. This paper reports on the state of a research initiative into human-intelligence/machine-intelligence (HI-MI) decision governance. The goal of this research is to develop an HI-MI decision governance body of knowledge as the systems policy basis for 21<sup>st</sup> century engineering systems management of HI-MI organizations. First, this paper lays a foundation of the basic HI-MI domains and their current state of research and knowledge. Next, the paper discusses the HI-MI intelligence implications for

engineering systems management in the 21<sup>st</sup> century. Finally, the paper describes the current state of this research initiative into HI-MI decision governance.

### Historical Overview of Human-Machine Intelligence Research

Acquisition or creation of human knowledge has always been directed toward explaining physical existence and improving human existence either tacitly or explicitly. In order to improve their existence, humans must apply gained knowledge through a decision process, either trial-and-error or logically designed, directed toward achieving an improvement goal. Thus, as illustrated in Exhibit 1, 21<sup>st</sup> century human and artificial intelligence research can be summarized along two domains: the intelligence domain and the decision source domain.

**Exhibit 1.** Domains of human and artificial intelligence.



#### Knowledge Engineering/Management (HI ⇒ HI)

Although some form of knowledge management has existed since humans first recorded paintings on cave walls or glyphs in stone tablets, modern knowledge engineering and management emerged in the late 1980s and early 1990s in the information and communication technology (ITC) domain. In its earliest ITC perspective, knowledge was considered to be in written form or existing in databases, e-mails, online libraries, etc. This has evolved into today's cloud based knowledge tools, blogs, discussion forums, social media, and wikis. It was quickly realized, however, that the information technology (IT) perspective alone was insufficient for capturing, encoding, and managing organizational knowledge. Three additional perspectives evolved in parallel with the IT perspective.

One perspective has been on how individuals create and share knowledge with a focus on building educational and knowledge sharing capabilities. Everett Rogers' (1962) work on diffusion of innovations contributed to understanding of how knowledge is created and diffused in social systems. Thomas Allen's (1977) work on evolved communications systems in science and engineering contributed to the understanding of the effects of informal and formal organizational structures on knowledge creation and dissemination. Peter Senge (1990) was one of the early researchers to focus on the cultural change required to create learning organizations. Argyris (1995) focused on how organizations work, evolve and learn. Nonaka and Takeuchi (1995) and Von Kroch, Ichijo, and Nonaka (2000) researched the dynamics of knowledge creation in business organizations. Argote, Miron-Spektor, Wang, and many other researchers continue work on understanding organizational learning.

Another focus has been on capturing and utilizing knowledge to improve enterprise effectiveness. Davenport and Prusak (1998) sought to explain how organizations generate, codify, transfer, and manage knowledge. Peter Drucker (2001) stressed the importance of organizational information and explicit knowledge as a valuable resource for improving competitive advantage. Dorothy Leonard (2005, 2014) has made multiple contributions to understanding creativity, innovation, and knowledge creation and management.

The final focus has been on exploiting IT to enhance enterprise economic value. Paul Strassman (1985, 1990, and 2007) was one of the early researchers that sought to answer the question of the economic value of information systems. Lesser and Prusak (2003) examined management methods for deriving tangible business value from knowledge management. Extensive research continues today into the economic value of organizational knowledge in general and within specific private, governmental, and education sectors.

## **Artificial Intelligence (HI ⇒ MI)**

Nilsson (2010) provides the following general timeline of AI development.

### **1950s**

- Pattern recognition – By the mid 1950s, the ability to scan images and convert them to number arrays (later termed “pixels”) had already been developed. Most of the early work in pattern recognition dealt with developing computer code to process these two-dimensional arrays, typed pages or photographs, for character recognition.
- Human learning, cognition, and memory – Initial work in human cognition resulted in the development of the first neural networks.
- Statistical methods – The earliest work was in statistical classifiers and extraction of distinguishing features from aerial photographs.
- Heuristic programs – Initial work focused on solving simple mathematical and geometric problems, solving puzzles, and playing games.
- Semantic representation – The initial focus was on how to organize information in a manner similar to how it is stored in human memory such that the “meaning” of words could be reconstructed in a humanlike way.
- Natural language processing – Work was initiated into converting natural language into an appropriate memory model or into attaining some action appropriate to the input.

**1960s** – Technical and societal developments converged to build the infrastructure needed for the development of AI. Faster, more powerful computers were developed, and the first specialized computer languages needed for symbolic manipulation were developed. Military support provided the means for the establishment of the first AI laboratories. Out of these laboratories came the first work in “hand-eye” research, which integrated cameras with rudimentary electromechanical prosthetic robotic hands and arms to manipulate simple objects.

### **1970s**

- Computer vision – Work was initiated into understanding the three-dimensional properties of human vision by translating and filtering differences in two-dimensional arrays to find edges and vertices objects from two stereoscopically mounted cameras.
- Processing line drawings – This work focused on how to segment line drawings into geometric segments (circles, squares, triangles, etc.).
- Robotics – By the end of the 1970s, the “hand-eye” research of the 1960s had evolved into mobile robots and robotics capable of assembling simple objects.
- Knowledge representation – By the end of the 1970s, work in knowledge representation resulted in the development of situation calculus, logic programming, semantic networks, and scripts and frames that are the basis of today’s expert systems and worldwide web knowledge retrieval.

**1980s** – With the foundation laid in the 1970s, AI work in the 1980s turned toward application.

- Speech recognition and processing – Work advanced the processing of continuous streams of speech and interpreting meaning.
- Consulting systems – Work was initiated into computers aided instruction for maintenance workers.
- Expert systems – Rudimentary expert systems were developed initially with knowledge about chemistry, spectroscopy, drug interactions, bacterial infections, and mapping mineral deposits.
- Computer visions – Research in computer vision advanced from finding edges and vertices and identification of basic geometric shapes to extracting properties of scenes and modeling solids.
- Japan’s Fifth Generation Computer Project, the British Alvey Program, Europe’s ESPRIT Initiative, and America’s Microelectronics and Computer Technology Corporation – All these projects and initiatives set similar goals of goal was to creating computers capable of AI inferences from large data and knowledge bases and communicate using natural language.
- DARPA’s Strategic Computing Program – Three major applications were initiated: (1) Pilot’s Associate to assist an air combat commander. (2) Battle Management System to assist the commander-in-chief of the U.S. Pacific fleet in planning and monitoring the operation of approximately 300 ships. (3) Autonomous Land Vehicle to use autonomous vehicles in combat, logistics and supply, and search and rescue.

### **1990s to date**

- Representation and reasoning – Work advanced to add nonmonotonic and defeasible reasoning to the monotonic reasoning developed in the 1970s and 1980s.
- Qualitative reasoning – Application of qualitative physics and mathematics knowledge to produce approximate solutions, which, in turn, can be applied to plan and execute subsequent qualitative estimates.
- Semantic networks – Development of description logics using frame-based knowledge representation.

- Constraint satisfaction problems – Assignment of values to an object such that the joint set of values satisfy a set of constraints.
- Propositional logic problems – Solving logical representation problems in which none of the logical formulas contain variables.
- Representing text as variables – Converting query text into vectors and retrieving documents with the closest matching document vectors.
- Latent semantic analysis - vector-based schemes for capturing meaning from text vectors.
- Causal Bayesian networks – Seeks to quantify the certainty of relationships among categorical factors and continuous variables by *a priori* probabilities and update that certainty based on natural and intervention state transitions.
- Machine learning – (1) Memory-based learning involves developing data reduction algorithms at the time data is retrieved to identify associations among data sets. (2) Case-based reasoning used pre-existing cases to sort through, analyze, interpret, and solve new cases. (3) Decision trees are constructed by extracting information from large databases to identify features in the data. (4) Reinforcement learning which of a large possible set of actions or decision should be executed to achieve an end result. (5) Deep {hierarchical} learning is based on algorithms that attempt to model high level abstractions, complex architectural structures, and non-linearities in data.
- Natural language processing – (1) Grammars and parsing algorithms that analyze natural language sentences and accept only legal word strings for meaning extraction. (2) Statistical natural language processing used to extract probabilistic meaning from ill-formed sentences.
- Computer vision – Purposive active vision to provide information needed for motor control.
- Cognitive system architectures – Parallel architectures designed to approximate intelligent problem solving behavior.

#### **Machine-to-Machine M2M Intelligence (MI $\Rightarrow$ MI)**

Machine-to-Machine (M2M) is the collection of technologies that enable “smart” sensors, actuators embedded processors, computers, and mobile devices to communicate with one another, take measurements, exchange data and information, and make decisions with and without human intervention.

M2M has existed in different forms since the advent of computer networking. The first working computer-to-computer network, constructed in the 1950s, was comprised of communicating computers that processed data from the Semi-Automatic Ground Environment (SAGE) radar system. In 1968, Theodore Paraskevakos combined computers and telephone systems to create the first caller line identification system. In 1977, Paraskevakos formed Metrotek, Inc. to develop and produce commercial remote meter reading and load management equipment, which eventually led to the smart meter and today’s concept of the smart grid. The hardware, software, and network (wired and wireless) components of M2M have been developed and produced by a relatively small group of manufacturers. Today, the primary applications of M2M include utilities monitoring systems, building monitoring systems, wireless digital billboards, industrial and petroleum control networks, and precision agriculture. Additionally, M2M is beginning to work its way into the HITECH (Health Information Technology for Economic and Clinical Health) system to permit exchange of patient data among hospitals, clinics, and doctors.

In the next few decades, M2M technology is forecast to evolve into the Internet of Things, including the Smart Grid, with the global adoption of the Internet Protocol Version 6 (IPv6) which expands the machine readable identifiers needed to make the Internet of Things a reality. Another evolving area of M2M technology is swarm intelligence (SI). Swarm intelligence, which evolved out of biological systems, is the collective behavior of decentralized, self-organized systems, natural or artificial. SI systems consist of a population of simple, lower intelligence agents, currently robots or drones, that interact locally with one another and their environment. Following simple rules, the agents jointly interact leading to the emergence of higher-order, global intelligent behavior toward the accomplishment of a mission

#### **Machine-to-Human M2H Intelligence (MI $\Rightarrow$ HI)**

To date, the human-machine interaction has been limited to human-computer or human-robot interaction. The Association for Computing Machinery (2015) defines human-computer interaction as "a discipline concerned with the design, evaluation and implementation of interactive computing systems for human use and with the study of major phenomena surrounding them." Research in the field of human-computer interaction is in:

- Designing optimized computer interfaces to achieve stated properties.
- Designing for ease of learning and use of computer applications.

- Developing methods for implementing human-computer interfaces to optimize human interaction with computers.
- Developing methods for evaluating and comparing interfaces for efficiency of use.
- Studying human computer use and its implications for organizational and broader sociocultural implications.
- Developing a body of knowledge of human computer use with conceptual frameworks for human-computer interfaces.

Loosely defined, human-robot interaction is just that; the study of interactions between humans and robots. There is no professional society devoted to the study of human-robot interactions; rather, there are a number of individual and laboratory research initiatives in diverse fields such as human-computer interaction, artificial intelligence, robotics, natural language understanding, design, and social sciences. The Association for Computing Machinery and the Institute of Electrical and Electronics Engineers hold an annual conference on human-robot interaction. The U.S. Navy's Center for Applied Research in Artificial Intelligence has a Cognitive Robotics and Human Robot Interaction research initiative (2015). The Center works under the hypothesis that

... robots and autonomous systems that use human-like representations, strategies, and knowledge will enable better collaboration and interaction with the people who use them. Similar representations and reasoning mechanisms make it easier for people to work with these autonomous systems. An autonomous system must be able to explain its decisions in a way that people understand, which should lead to better trust and acceptance of the system. If an autonomous system can predict a person's needs, even in the very short term, it can prepare for it and act appropriately.

From this hypothesis, the Center has two primary scientific goals:

- To understand the embodied nature of cognition: how people work in the physical world.
- To improve human robot interaction by high fidelity models of individuals so that we can provide some assistance to them.

Like the rest of the AI community, the U.S. Navy's human-robot interaction research initiative assumes that human intelligence can be so precisely described and modeled that it can be completely simulated in AI. Thus, humans become discontinuities to be accommodated by AI robotics.

### **HI-MI Intelligence Implications for Engineering Systems Management**

Given unquestioned overarching assumption of its ability to completely model human cognition by the AI community, there has been no research into cooperative human-machine decision making. Further, given the AI community's continuing failure to come close to this goal let alone attain artificial intelligence as specified by the Turing test (ability of a computer or machine to exhibit intelligent behavior equivalent to, or indistinguishable from, that of a human), there has been no research into the human-intelligence/machine-intelligence decision governance needed to integrate and manage cooperative human and machine decision making in the organization. The implication for the Engineering Management discipline is immediate. As with the U.S. medical community's HITECH implementation, there will be major failures and unintended consequences as machine intelligence is integrated more and more into engineering and engineering-supported organizations.

### **Research into Human-Intelligence/Machine-Intelligence HI-MI Decision Governance**

To address this gap, research was initiated toward developing an HI-MI decision governance framework and body of knowledge with the ultimate goal of building an Engineering Management HI-MI decision governance discipline. This research is built on two premises and one proposal.

**Premise 1:** Artificial intelligence can only asymptotically approach general human cognitive intelligence.

**Premise 2:** Due to limited human capacity to process information, artificial machine intelligence can be developed to approach, equal, and potentially outperform humans in some domain specific decision tasks.

**Proposal:** Designed integration of human intelligence and machine intelligence within an HI-MI decision governance framework can produce robust systemic mission accomplishment; that is, achievement of a specified set of systemic objectives while minimizing the probability of serious or fatal errors under widely varying environments with potentially nonlinear, discontinuous risky and uncertain events.

Premise 1 is self-evident, because of the limited cognitive capacity of humans to understand their own tacit and explicit cognitive processes. Further, the human brain is elastic. As humans understand more about their own cognitive capacity, the human brain will create new tacit knowledge about its own cognition processes (i.e. as problems are solved, new unknowns will be identified) at a velocity ahead of that which humans can achieve to capture and convert the new tacit knowledge into actionable explicit knowledge. Premise 2 is self-evident because it has been demonstrated already that in existing domain specific tasks, machine intelligence does outperform human decision making and problem solving capacity. The resulting proposal establishes the general research framework; that is, developing a general theory and body of knowledge of HI-MI decision governance with a focus on systemic mission accomplishment within widely varying risky and uncertain environments. For this research, the proposal differentiates machine intelligence from general artificial intelligence and delimits the definition of human-intelligence/machine intelligence.

**Definition:** Machine intelligence is the artificially intelligent, domain specific decisions and actions required to accomplish a specified systemic mission.

**Definition:** Under imperfect human intelligence and imperfect machine intelligence, joint human-intelligence/machine-intelligence is the optimal, bounded set of systemic, domain-specific decisions and actions required for a system of human-machine agents to accomplish a specified systemic mission.

The HI-MI definition allows for a joint set of human and machine decisions and actions toward a specified mission under evolving states of human-intelligence/machine-intelligence. The optimal set of systemic decisions and actions toward mission objectives is bounded only by the current state of knowledge about human-intelligence/machine-intelligence within the domain specified mission context. As the state of knowledge increases, the definition of optimal decisions and actions can be refined to reflect the reduced risk and uncertainty in outcomes.

Exhibit 2 sets forth the integrative approach toward developing the general HI-MI decision governance theory and body of knowledge. The general research approach seeks to integrate existing socio-technical systems knowledge with decision theory and AI declarative and procedural knowledge into a human-intelligence/machine-intelligence systems theoretical framework and body of knowledge and then validate it through causal modeling of specific organizational decision instances. The approach integrates existing systems, knowledge, and data governance bodies of knowledge to build the viable system policy level five (Beer, 1994) HI-MI decision governance framework. The level four external environment monitoring and future planning framework will be built from general systems theory and decision theory prediction and planning. The level three self-organization and regulation will be built from socio-technical and organizational theory. The level two anti-oscillatory framework will be built on systems cybernetic theory. Current state AI and human-machine interactions knowledge will form the basis for the level one operations framework. The human and machine agents will be integrated as level one, recursively intelligent entities whose joint decisions and actions are directed by the levels two through five policy and cybernetic control structure toward systemic mission accomplishment.

**Exhibit 2.** Integrative approach to HI-MI decision governance theory and body of knowledge.

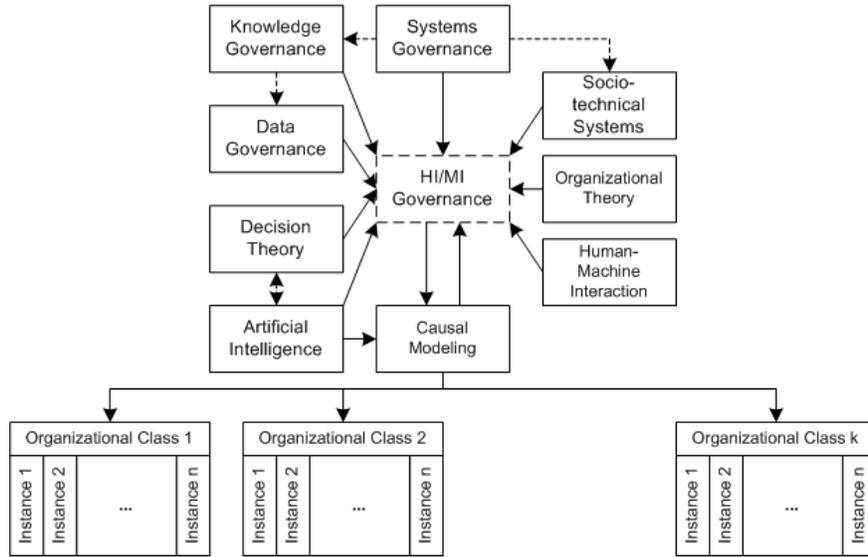
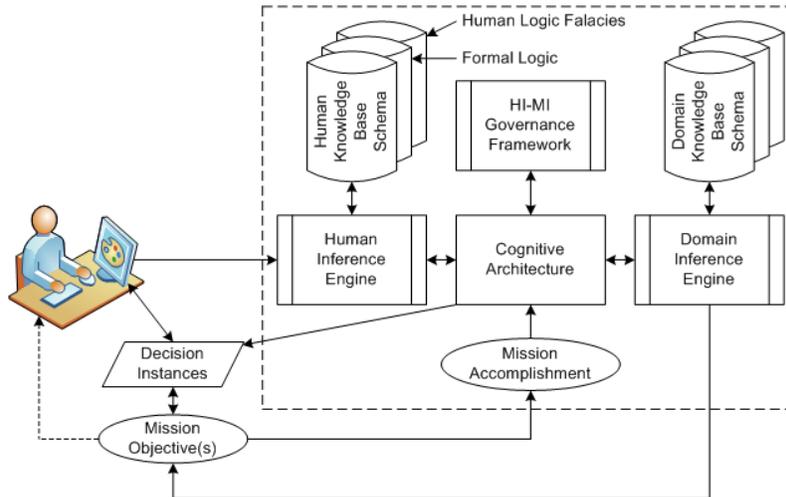


Exhibit 3 sets forth the HI-MI decision governance architecture for implementing, testing, and validating the theoretical framework and body of knowledge. The viable system decision governance framework will be encoded in a collection of knowledge representation ontologies and instantiated in a cognitive architecture. Existing engineering, mathematical, and physics ontologies will form the basis of the basic domain knowledge and domain specific ontologies will be imported, where they exist, or built to reflect currently accepted optimal decision processes in a respective domain. The human knowledge base will be built on three ontologies: (1) first order propositional, probabilistic, and abductive logic, (2) human logic fallacies, and (3) individual human subject decision making pathologies encoded from the human inference engine as deviations from first order logic or following logic fallacies. The cognitive architecture will be optimized for each specific organizational decision instance to select the human-intelligence/machine-intelligence decision inputs that jointly minimize deviations from a vector of mission objectives. A knowledge base of HI-MI decision processes will be built for each organization class, and organizational decision process knowledge bases will be concept mined for common HI-MI decision governance rules.

**Exhibit 3.** HI-MI decision governance architecture.



## Conclusions and Future Research

Current research is progressing toward building the HI-MI governance architecture illustrated in Exhibit 3 and building the systems, knowledge, and data governance ontologies that will form the basis of the viable system policy level five HI-MI decision governance framework. When the HI-MI governance architecture has been completed and tested and the HI-MI decision governance framework encoded and validated, work will proceed to building the domain and human knowledge bases.

## References

- Allen, T. (1977). *Managing the Flow of Technology*. Cambridge, MA: MIT Press.
- Argyris, C. (1995). *On Organizational Learning*. New York: Wiley.
- Association for Computing Machinery. (2015). *ACM SIGCHI Curricula for Human-Computer Interaction*. Retrieved from [http://old.sigchi.org/cdg/cdg2.html#2\\_1](http://old.sigchi.org/cdg/cdg2.html#2_1).
- Beer, S. (1994). *Brain of the Firm*. New York: Wiley.
- Davenport, T. and Prusak, L. (1998). *Working Knowledge: How Organizations Manage What They Know*. Boston: Harvard Business School Press.
- Drucker, P. (2001). *Management Challenges for the 21<sup>st</sup> Century*. New York: HarperCollins Publishers.
- Leonard, D. and Swap, W. (2005). *Deep Smarts: How to Cultivate and Transfer Enduring Business Wisdom*. Boston: Harvard Business Review Press.
- Leonard, D., Swap, W., and Barton, G. (2014). *Critical Knowledge Transfer: Tools for Managing Your Company's Deep Smarts*. Boston: Harvard Business Review Press.
- Lesser, E. and Prusak, L. (2003). *Creating Value with Knowledge: Insights from the IBM Institute for Business Value*. New York: The Oxford University Press.
- Nilsson, N. (2010). *The Quest for Artificial Intelligence*. New York: Cambridge University Press.
- Nonaka, I. and Takeuchi, H. (1995). *The Knowledge-Creating Company*. New York: The Oxford University Press.
- Rogers, E. (1962). *Diffusion of Innovations*. Glencoe: Free Press.
- Senge, P. (1990). *The Fifth Discipline: The Art & Practice of The Learning Organization*. New York: Doubleday.
- Strassmann, P. (1985). *Information Payoff: The Transformation of Work in the Electronic Age*. New York: Free Press.
- Strassmann, P. (1990). *Business Value of Computers: An Executive's Guide*. New Canann, CT: Information Economics Press.
- Strassmann, P. (2007). *The Economics of Corporate Information Systems: Measuring Information Payoffs*. New Cannan, CT: Information Economics Press.
- Sreenivasin, H. (May 8, 2015). "How smart is today's artificial intelligence?" Retrieved from PBS NEWSHOUR: <http://www.pbs.org/newshour/bb/smart-todays-artificial-intelligence/>.
- United States Navy. (2015). *Cognitive Robotics and Human Robot Interaction*. Retrieved from <http://www.nrl.navy.mil/itd/aic/content/cognitive-robotics-and-human-robot-interaction>.
- Von Krogh, G., Ichiho, K., and Nonaka, I. (2000). *Enabling Knowledge Creation*. New York: The Oxford University Press.
- Wachter, R. (2015). *The Digital Doctor*. New York: McGraw Hill.

## About the Author

**T. Steven Cotter** is a Lecturer with the Engineering Management and Systems Engineering department at Old Dominion University. He earned a Ph.D. in Engineering Management and Systems Engineering from Old Dominion University, a Master of Science in Engineering Management with a concentration in quality/reliability engineering from the University of Massachusetts at Amherst, a Master of Business Administration with a concentration in finance and a Bachelor of Science both from the University of South Carolina, and a diploma in Electronic Technology from Graff Area Vocational and Technical School (now Ozarks Technical Community College). He is a certified Quality Engineer and Reliability Engineer with the American Society for Quality. His research interests are in engineering analytics design, human-machine intelligent socio-technical organizations, quality systems design, and statistical engineering.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.